




# Building Responsive Data Streams for Modern Healthcare Systems

**Pierre Moreau**

Asst Professor

## Abstract

Data-driven disruption in healthcare enables advanced analytical capabilities for near-real-time monitoring, prediction, and prevention. However, existing interoperability standards remain inadequate for large-scale operational systems. An innovative event-driven architectural pattern addresses the requirements for near-real-time healthcare data streaming and corresponding data models. Key streams in acute care and telemetry from medical devices illustrate the application of these concepts.

Research demonstrates that many common event-driven architectural patterns well-known in business domains apply equally to near-real-time infrastructures in healthcare. Data-driven technologies are rapidly disrupting business and government sectors. Purpose-built data pipelines collect, curate, and stage data for near-real-time operational insights into customer and environmental factors. Emerging analytical techniques such as complex event processing, predictive analytics, and machine learning rely on these data streams for service innovation, business performance and resilience, fraud detection, and automated response. Similar capabilities are desirable in healthcare, whether for monitoring ICU patients, predicting hospital readmissions, preventing clinical deterioration, detecting adverse events, optimizing resource allocation, or curtailing the spread of disease. Nevertheless, the healthcare sector has not yet realized these advanced operational insights. The primary reason is the lack of sufficiently low-latency, scalable healthcare data streams. Drivers of change include research demonstrating suboptimal patient outcomes resulting from failure to act on existing evidence, the infectious nature of COVID-19, concerns regarding healthcare system capacity, and a growing risk of influenza-H1N1 co-infection.


**Keywords :** Event Streaming, Real-Time Data Processing, Healthcare Interoperability, HL7 / FHIR Integration, Message Queues, Apache Kafka, Event Brokers, Data Pipelines, Microservices Architecture, Clinical Data Streams, Asynchronous Communication, Data Ingestion Frameworks, Patient Monitoring Systems, Low-Latency Processing, Healthcare Analytics Streams.

## 1. Introduction

The continuous, growing body of data from near-real-time healthcare is sourced from a multitude of entities—among them, sensors, medical devices, devices in support of active patient treatment, tools in support of diagnostics, surveillance cameras/pooled data from general settings for COVID-19 purposes (akin to control rooms), pharmacy, aided/loss events—along with messages from telemetry from surgical rooms to applied scientific simulation. These bodies of information are comprised of data, data about data (metadata), timing (temporal alignment), and contextual

attributes, along with methods for continually detecting key body, disease, and pathology conditions. The sources constitute diverse, distributed data networks of event sources, and they must deal with complexities that make their continuous processing and analytics challenging, particularly with regard to data velocity and volume. The produced data velocity requires processing and analytics of an event stream at a frequency that can capture the desired underlying reality. Data volume requires careful trade-off evaluations to reduce data storage and communication costs, while ensuring that business objectives continue to be achieved. Naturally.

The key to the solution must be an event-driven architecture (EDA) that provides a near-real-time stream of appropriately merged and transformed events. Central to the event-driven architecture for streaming near-real-time healthcare data is the use of publish-and-subscribe and event-sourcing patterns to support complex event processing, continuous analytics, and near-real-time data lakes feeding ML and AI models that require high event-ingestion rates. Central, supporting quality-assurance processes provide information about certain conditions not covered in the original event, supporting decision-making relative to those conditions.



**Fig 1: Real Time Event Streaming In Healthcare**

### 1.1. Research design

This research analysis proposes a security, privacy, and compliance-sensitive near-real-time healthcare event-driven architecture (EDA) using negative and positive security models. To achieve near-real-time, compliance- and security-sensitive health data streaming, the core requirements of an EDA are expanded. Important near-real-time healthcare streaming requirements are high throughput, low latency, failover reliability, fallback reliability, data protection, privacy control, and auditability. Support patterns validated in related research are identified, and five classes of such support patterns are evaluated for suitability.

A research direction to provide support for three streaming classes—medical-device-generated telemetry, population and postacute care healthstreams, and low-latency acute-care

sensor data—is outlined. The direction includes a future direction along an open-source front, with details for one streaming class involving a non-proprietary, low-cost blood glucose monitoring device. Individual acute-care remission deployment risks are low and can be contained within acute-care settings with reliable latency. Remaining risks warrant research and pilot partnerships with manufacturers of policy-governed, policy-agnostic, and community-cloud-enabled medical devices and software-as-infrastructure platforms to build cloud-based telemetry pipelines.

### Equation 1: End-to-End Latency

#### Step 1: Break the pipeline into stages

For one healthcare event, the delay comes from:

- sensing/data generation delay
- transmission delay from producer to broker
- broker queuing delay
- processing delay
- delivery delay to consumer/dashboard/analytics

Let these be:

- $L_g$  = generation latency
- $L_t$  = transmission latency
- $L_q$  = queueing latency
- $L_p$  = processing latency
- $L_d$  = delivery latency

#### Step 2: Total latency is the sum of stage delays

A message must pass sequentially through these stages, so the total is additive:

$$L_{total} = L_g + L_t + L_q + L_p + L_d$$

## 1.2. Background and Significance

Near-real-time healthcare data streaming has the potential to fundamentally transform health care delivery by enabling actionable intelligence for providers and by facilitating a broader array of products and services. Event-driven architecture (EDA) patterns that can support near-real-time healthcare data streaming are examined from a security, interoperability, and reliability perspective. EDA enables the application of real-time analytics to streaming health data in addition to delivering data to consumers for retrospective analysis. As near-real-time healthcare data streaming matures, there is a clear need for structured, reusable health data models and semantics to complement the existing focus on event processing and real-time analytics.


The case for healthcare event-driven architecture has been made in several previous studies for domains where latency and real-time analytic capabilities are critical (acute care, perioperative services, telemetry listening) and where there is a confluence of multiple data sources (post-acute care, population health, etc.). Events serve as the building blocks for streams, and streams serve as the building blocks for messages. The design patterns and principles of structured event-driven architecture apply equally to near-real-time healthcare data streaming and to near-real-time healthcare data analytics. Event processing, complex event processing, and streaming health data all benefit from the application of a publish-subscribe model, and the case for such a model extends beyond event-driven architectures with true real-time requirements. Recent engagements in the greater Boston area focused on healthcare cloud platforms for near-real-time analytics and streaming health data.

## 2. Security, Privacy, and Compliance Considerations


Effective access control mechanisms and data-in-transit encryption minimize unauthorized use and reduce the risk of data loss. However, near-real-time data streaming often occurs over public cloud services that, though highly secure,

are not governed by the same data-security framework as private hospital systems. When data are sent from hospitals “over the wire” to be processed and stored off-site, the risk of exposure remains. Pseudonymization or encryption of medical data before leaving the hospital’s local environment may therefore be necessary to comply with data minimization principles and cope with the sensitive nature of healthcare data. Data-minimization requirements may also be imposed by connecting organizations, and so a balance must be struck between the need for protection and the need for useful analysis.

Healthcare data comprise a rich source of insights for a variety of stakeholders throughout the healthcare ecosystem, some of which are commercial for-profit organizations, such as insurance providers. These use cases are thus typically much less stringent than the safeguards required to enable direct patient data to be shared with other clinical organizations without explicit consent. Nevertheless, supporting GDPR and HIPAA regulatory frameworks requires the implementation of audit trails, including cryptographic monitoring techniques, to proof the processed data against privacy, confidentiality, and consent aspects and to establish provenance tracking.



**Fig 2: Healthcare data security and privacy in Data Warehouse architectures**



### 2.1. Data Security in Transit and at Rest

Sensitive data in transit is particularly vulnerable to unauthorized access and modification during transmission. Because numerous standards in the healthcare domain mandate that data be encrypted both in transit and at rest, implementing cryptographic measures is a critical design consideration in a near-real-time data-sharing architecture. Special attention must be paid to ensure that sensitive data is encrypted in such a way that it remains in a secure state even during processing. For instance, a protocol like Intel's Software Guard Extensions enables an application to host sensitive data in memory in an unencrypted form only within an execution environment guaranteed to be inaccessible from other processes running on the same machine.

However, encryption in transit lacks protection against endpoint vulnerabilities. Many European and national directives have therefore stipulated the restriction of sensitive data exposure during processing by promoting the principles of data minimization and purpose specification. The principle of least privilege also requires that only authorized personnel have access to sensitive data items.

The explicit requirement for limited access to specific groups of users is inherent in the implementation of regulatory measures such as the Health Insurance Portability and Accountability Act, the Food and Drug Administration Data and Safety Monitoring Board guidelines, and European Union legislation. Security and privacy compliance in health data streams is also a concern for industry standards like the Fast Healthcare Interoperability Resources. Audit trails complement access restrictions by generating tamper-proof logs of all accesses to sensitive data in health data systems.

| Concept | Definition                    | Role in Healthcare            |
|---------|-------------------------------|-------------------------------|
| Event   | Change of state or occurrence | Patient vitals change, alerts |
| Stream  | Ordered sequence of events    | Continuous telemetry data     |

| Concept        | Definition                    | Role in Healthcare            |
|----------------|-------------------------------|-------------------------------|
| Message        | Data unit carrying event info | Transmits clinical updates    |
| Event Source   | Producer of events            | Sensors, EHR systems          |
| Event Consumer | Subscriber to events          | Dashboards, analytics systems |

**Table 1: Core Concepts of Event-Driven Architecture**

### 2.2. Access Control and Data Minimization

Access controls based on the principle of least privilege limit health system and third-party access to only the data, resources, and services required for specific health care delivery and business functions. Security and privacy risks are minimized by restricting the use of health data for marketing, selling, or other non-health care related operations. In addition to technical access controls, carefully designed user authentication and training programs are necessary to reduce the risk of human error. Even with these security measures, health data may be inadvertently or maliciously disclosed or accessed. Data-minimization principles reduce the risk of sensitive data being misused or inappropriately disclosed. Within the health-care context, these principles include restricting the display of identified data to direct care providers, protecting the identities of patients undergoing sensitive medical treatments or tests, and limiting access to health data retained for business or research purposes.

In clinical settings, these principles are supported by patient consent. Sensitive data are identified within the shared health data model, a data-sharing standard proposal inspired by the Fair Information Practice principles. The principle of purpose specification states that health data should only be used for the purpose expressly consented to by the patient; for example, marketing or research-related use. Whenever possible, the purpose, data being used, and the period

covered for the deployment of the health data should be specified. Further, public communicating processes should be established and patients should be kept informed continuously. Since patients have various sensitivities, further privacy protection and purpose specification are necessary for patients undergoing drug-abuse treatment or sexual-transmitted-disease diagnosis and treatment.

## Equation 2: Throughput of the Streaming System

### Step 1: Define throughput

Throughput means the number of events processed per unit time.

If:

- $N$  = number of events processed
- $T$  = total time interval

then throughput is

$$\theta = \frac{N}{T}$$

### Step 2: Byte-based formulation

If each event has average size  $s$  bytes, then data throughput is

$$\theta_b = \frac{Ns}{T}$$

### Step 3: Relation to service rate

If the system can process one event in average service time  $t_s$ , then service rate is

$$\mu = \frac{1}{t_s}$$

For  $k$  identical parallel processors, maximum ideal throughput becomes

$$\theta_{max} = k\mu = \frac{k}{t_s}$$

**2.3. Regulatory Frameworks and Audit Trails** Near-real-time healthcare data streams operate within a complex security and regulatory landscape. By design, these streams may carry sensitive health information between diverse providers and organizations, exposing them to an elevated risk of re-identification and other security breach scenarios. Consequently, data minimization techniques are essential for reducing that risk. Often, however, the operators of such streams have to provide guarantees for regulatory data governance. This often requires maintaining a complete audit trail of all events contained within the stream, along with evidence or copyright-like certificates of proper key use. Certain regulatory compliance frameworks, such as HIPAA in the US or GDPR in the EU, require such an audit trail for certain types of sensitive healthcare data in the event of a security or privacy breach.

Audit trails can be built on top of an event sourcing pattern, where each received event is produced as the serialized health event, for regulatory and governance purposes. All pull subscribers of the data lake in replay-mode and any data pipeline performing data enrichment can verify that they indeed operated properly through ownership of the original event and by private data minimization techniques on the sensitive attributes. For the streaming systems that act as the main event source, idempotent processing and storage are required to avoid spurious re-publication of events, since such re-publication occurs through the publish-subscribe backend.

## 3. Fundamentals of Event-Driven Architecture in Healthcare


After defining events, streams, and messages in the context of an Event-Driven Architecture (EDA) for streaming healthcare data, this section illustrates the functional

components of an EDA in healthcare and examines how data provenance, quality, and governance are assured in the architecture.

### Events, streams, and messages

An event is an occurrence or change-of-state in a system that is potentially of interest. A stream is an infinite sequence of ordered events that happen over time. A message is a transactional unit of data transmission used to convey intent or effect change.

An EDA uses events and streams as core abstractions. An event represents a meaningful change in the state of the system, the occurrence of an action, or a piece of information that may be relevant to other participating components. Each stream consists of a sequence of events of the same type expressed in temporal order and is continuously updated as new events occur. An event may also be represented as a message that is published on a stream. In the future, such published messages enable notifying other components interested in the change without direct coupling between components. —inspired by Event-Driven Architecture for Near-Real-Time Healthcare Data Streaming.



**Fig 3: Event-Driven Architecture**

### 3.1. Events, Streams, and Messages

Healthcare EDAs produce, consume, and process event data. Events capture changes of state originating in business processes. Healthcare events are typically state changes in an

individual patient, but they may also reflect group-level changes, such as the onset of the COVID-19 pandemic; administrative changes in guidelines, standards, insurance coverage, and reimbursement; and demographics and social determinants of health.

An event stream is a high-fidelity template or schema defining a collection of events. It describes a business concept important for data-sharing and analytics use cases. Event-stream producers publish individual events. The relationship between events and event streams is similar to that of a message queue and its message schema. An event stream is a standardized template or schema defining a collection of individually published healthcare events. Stream-schema evolution is similar to that of database schemas, with a standard process for applying backward-compatible and breaking changes.

Event messages contain the domain-specific data about the state change, as well as meta-data about the event, such as its time-stamp and origin. Meta-data defines provenance, granularity, and expressed data quality. Events are discoverable at an attribute level based on the attributes defined in the corresponding stream schema. Schema evolution may include adding or removing individual attributes or changing data types. Event messages are “transactions” within the event stream. They can also be expressed in real-time Query Language (rQL) syntax in many EDA implementations. Querying at this level provides a powerful capability to query recent event history in different persistency stores, for use cases ranging from detection to detailed diagnosis.

### 3.2. Components of an EDA in Healthcare

Eight components form the core of a well-designed event-driven architecture for near-real-time data streaming: a data lake, a data ingestion pipeline, a streaming platform, an event-function store, a complex-event processor, real-time dashboards, and an event-streaming producer. The data lake serves as an analytical repository that integrates all available clinical data. It is complemented by the ingestion pipeline,

which assimilates historical data from data sources at rest, such as a data warehouse, and scrubs it for quality before depositing it into the lake.

The data lake connects to a streaming platform that manages data in motion. The platform offers health systems and their third-party partners a publish-subscribe mechanism to share monitoring events with minimal latency and provides the foundation for a managed service for clinical event processing. Monitors or web services observably commit events to the streaming platform. These events are typically patient-centric or population-centric content, such as SOAP (Simple Object Access Protocol) messages containing the full FHIR encoding of a patient-centered clinical encounter.

### Equation 3: Queue Growth and Backpressure Condition

#### Step 1: Define rates

Let:

- $\lambda$  = arrival rate of events
- $\mu$  = processing/service rate
- $Q(t)$  = queue length at time  $t$

#### Step 2: Net queue growth

Over a small interval  $\Delta t$ :

- events arriving =  $\lambda \Delta t$
- events served =  $\mu \Delta t$

So the change in queue is

$$\Delta Q = (\lambda - \mu) \Delta t$$

#### Step 3: Differential form

Divide by  $\Delta t$  and let  $\Delta t \rightarrow 0$ :

$$\frac{dQ}{dt} = \lambda - \mu$$

#### Step 4: Backpressure condition

If arrivals exceed service:

$$\lambda > \mu$$

then

$$\frac{dQ}{dt} > 0$$

### 3.3. Data Provenance, Quality, and Governance


Data provenance, quality, and governance are paramount considerations for a healthcare data streaming architecture. Interoperability requires that data is meaningful when interfacing with other systems, yet limiting access to data that can change state or is still evolving can reduce latency. Security regulations, such as the U.S. Health Insurance Portability and Accountability Act, mandate that healthcare organizations actively protect the data streams they create. Government and industry groups provide supplemental assistance by creating specialized standards and protocols for healthcare data transfer.

Provenance refers to documentation of the origin of the data and prior stages of its existence. Provenance data for a healthcare data stream must include the authoritative source of data generation for every event record in the stream, something akin to a forensics chain of custody. Provenance is essential for the trust models driving data-sharing agreements. Privacy regulations typically require audit trails to track user access and give patients the right to a report of which entities have accessed their data. For data lakes that ingest streams from multiple sources, third-party health data industry standards such as STRIDE may assist in these compliance and provenance needs.

## 4. Near-Real-Time Data Streaming Requirements in Healthcare

Most therapeutic and diagnostic concepts in healthcare are time-sensitive; therefore, stream processing is often preferred over batch processing for implementing AI and ML models. However, healthcare is also complex. Advanced-analytic systems often require copious amounts of diverse data from multiple sources and sensors, not all of which generate streams or provide such data reliably. The healthcare setting—as well as other contexts characterized by imperfect or unreliable data—introduces additional considerations beyond throughput, latency, and reliability. These considerations include security, privacy, architectural design, and external-facing interfaces.

Healthcare organizations must protect their patients' data, prevent information systems from becoming conduits for attacks, and adhere to the requirements codified in HIPAA while enabling use-cases that generally rely on data moving from an event producer to a consumer. "Where to put it/to whom to give it" becomes a corresponding high-level design concern, compounded by the complexity of healthcare services—actors, events, domains, and departments inevitably interact—and potentially fulfilled externally by standards such as HL7 FHIR, XDS, and Integrating the Healthcare Enterprise (IHE). Visualizing the outflow of information as a series of data streams or pipelines for further use will also appear to facilitate architectural comprehension and "make sense" to developers.



**Fig 4: Real-Time Data Streaming Architecture**

#### 4.1. Latency, Throughput, and Reliability

Latency, the time between input and output, is a key design constraint in realtime, interactive systems. For healthcare, the requirement is common but has unique properties. Latency encompasses both data generation and processing. It can be defined as the maximum allowable time for a complete cycle: the full round-trip time of a typical message—a heartbeat for wearable devices, a few minutes for other critical data, an hour for most of the rest.

Healthcare events require more than low latency. For obvious reasons, there can be no loss of information; consumers must receive all input messages in order; and additional constraints apply during periods of high load. Producing systems must not be overloaded, yet events must not accumulate undelivered for long. Although interruptions can normally be tolerated, a system lacks redundancy if both the producer and consumers are down at once. Fault-tolerance mechanisms should ensure that the required message guarantees, including temporal ordering, are satisfied throughout normal operation and during recoveries.

#### 4.2. Security, Privacy, and Compliance

Both privacy and security procedures must be included in the collection and sharing of health data for robust and trustworthy systems. Compliance and audit capabilities must also be embedded into the infrastructure to allow organizations to meet regulatory frameworks such as General Data Protection Regulation, Health Insurance Portability and Accountability Act, Federal Information Security Management Act, Sarbanes–Oxley Act, and Personal Information Protection and Electronic Documents Act. The audit procedures help detect potential breaches and offer visibility that permits organizations to demonstrate compliance before auditors. The design of the application must thus include elements of access control, logging, retraining, and encryption.

For access control, only authorized endpoints, such as analysis and storage facilities, should be allowed access to the sensitive data. Such facilities should also collect only

those data essential for their functions. For logging, health records within a cloud service must be monitored: on whom they were accessed, when they were accessed, what operations were performed, and the data usage. Data at rest may be protected by strong encryption, which renders the data unreadable without the important keys stored safe in a secure element. When an important key is compromised, automatic key rotation will replace a compromised key in a timely manner. These elements help maintain privacy during data sharing but do not ensure that sensors whose operations are intrusive are properly managed and do not violate the privacy and autonomy of patients.

| Component               | Description                | Function                                     |
|-------------------------|----------------------------|----------------------------------------------|
| Data Lake               | Central storage repository | Stores structured & unstructured health data |
| Ingestion Pipeline      | Data collection mechanism  | Cleans & loads historical data               |
| Streaming Platform      | Real-time data handler     | Manages data in motion                       |
| Event Store             | Event persistence system   | Maintains event history                      |
| Complex Event Processor | Analytics engine           | Detects patterns & triggers alerts           |
| Real-Time Dashboard     | Visualization layer        | Displays patient/clinical data               |
| Event Producer          | Data generator             | Devices, applications                        |
| Event Consumer          | Data receiver              | Clinical systems, analytics tools            |

**Table 2: Components of Healthcare Event-Driven Architecture**

### 4.3. Interoperability and Standards

In addition to being secure, reliable, and scalable, a near-real-time data streaming architecture must interoperate seamlessly with the many systems, devices, and services in a hospital or health system. Standards play an essential role in achieving this interoperability, but it is also important to examine query-based publish-subscribe systems or monitor-access patterns and not restrict the discussion to support for notifications.

No two systems trigger events in precisely the same way. Such variability at event producers complicates the design of event consumers—the systems that listen for, receive, and process the events. Data types and schemas can vary across systems and application domains, thus creating additional integration challenges. In data integration, both the interoperation of multiple potential realizations of a data element (data source systems) and the interoperation of all the consumers (data sink systems) must be considered. These requirements are often referred to as reciprocity and reciprocity of function. If shameful, public consumption is permissible for data integration, then support for shaming becomes an additional integration requirement. Such variations among systems can perhaps be mitigated through a central event bus design that recognizes and contextualizes these disparities. Nevertheless, for scalable, distributed event-based monitoring, some shorthand or shared semantics are essential.

The need for scalable, highly available systems, along with the complexity and heterogeneity of the monitoring infrastructure, greatly encourages the use of standard formats for the exchanged data. Such Open Systems Interfaces take on added importance as the scale of the architectures increases because well-defined interfaces reduce the overhead of the inevitable consulting that occurs during implementation. Fortunately, many Healthcare Systems already employ standard event formats in their monitoring functions (e.g., syslog for UNIX vendors and SNMP for network devices). Consequently, it is reasonable to expect that other classes of systems can also adopt standard event record formats, such as the Common Audit Event Format

promulgated by the Internet Engineering Task Force and the protocols of the Web Services Choreography Description Language Working Group.

#### Equation 4: Event Replay / Event Sourcing State Reconstruction

##### Step 1: Define state transition

Let:

- $S_0$  = initial state
- $e_1, e_2, \dots, e_n$  = ordered events
- $f(S, e)$  = transition function that updates state using event  $e$

After first event:

$$S_1 = f(S_0, e_1)$$

After second event:

$$S_2 = f(S_1, e_2) = f(f(S_0, e_1), e_2)$$

Continue similarly.

##### Step 2: Generalize after $n$ events

The final state is

$$S_n = f(f(\dots f(S_0, e_1), e_2) \dots, e_n)$$

##### Step 3: Operator notation

Define an event application operator  $\circ$ , then

$$S_n = S_0 \circ e_1 \circ e_2 \circ \dots \circ e_n$$

## 5. Architectural Patterns for Near-Real-Time Streaming


Streaming data demands a set of architectural patterns that address key operational requirements such as high throughput, low latency, elasticity, scalability, and reliability. The discussion focuses on five principal patterns: publish-subscribe and event sourcing, complex event processing and real-time analytics, and the ensemble of data lakes, data pipelines, and streaming platforms.

In healthcare, the publish-subscribe pattern is prevalent, where data producers publish events on specific topics and event consumers subscribe to the topics of interest. For example, a heart monitoring system may publish heart activity and heart failure-preventing events, while the attending physician may be interested only in the heart failure-preventing events. Stream processing systems (e.g., Kafka and RabbitMQ) provide the infrastructure to implement the publish-subscribe pattern. Two additional advantages of these systems are scalability and fault tolerance based on the replication of events across distributed systems.

The event sourcing pattern is an approach for modeling the system's state based on the series of events instead of the state of data. The current state can be reconstructed by replaying the previous events. This architecture pattern is beneficial when the overhead for reconstructing the present state is small with respect to reloading the state from a persistence layer. This pattern is also important in support of auditability (by following the events) and for enforcing data integrity and consistency.

Another architectural pattern ensures that a data stream is continuously monitored to detect specific conditions and to take proactive actions. For example, in an acute healthcare setting, respiratory patterns, heart activity, and medication administration can be used to identify patients who might develop new complications. The pattern addresses not only monitoring but also treatment. An event can trigger a query such as "Is there a bed available?" in a hospital ward and actuate the response: "Transfer patient X to ward Y."

Finally, the streaming analytics and complex event-processing patterns are increasingly supported by services from public cloud providers. Complex event processing techniques help detect patterns that span multiple data streams and can trigger alerts when specific situations arise. Cloud streaming analytics support serverless executions, making it possible to deploy a small piece of code that will be executed only when a specific analysis must be conducted.



**Fig 5: Near real-time solution – an architecture that works - Practical Real-time Data Processing**

### 5.1. Publish-Subscribe and Event Sourcing

A publish-subscribe architectural style, in which sources of events publish them to a common broker that delivers messages to interested subscribers, is often used in combination with event sourcing for near-real-time streams in healthcare. In event sourcing, state transitions of a data entity are persisted as a sequence of events rather than as just the final state. Each update or change to the data is captured as an event. Because events maintain the intermediate state of the resource, the prior states of the data can be reconstructed by reprocessing the events in sequence. Compliments have been noted with event sourcing and CQRS (command-query responsibility segregation) patterns, in which these two concerns of data update and query can be handled with a data model that is optimal for each. Simple event-processing components can publish changes to the data as event streams. These can be leveraged to provide CQRS data models for downstream system architecture.

Publish-subscribe, event sourcing, and CQRS patterns can be combined together in the delivery of telemetry streams from wearable, implanted, or consumed devices (e.g. wearables, ingestables) for acute care settings, as well as for chronic-care and longitudinal monitoring use cases. Event streams in these scenarios are often implemented as telegraf/grafana signalcloud, capable of also visualizing at scale the sensed values in chronological order. Publish-subscribe event streams also enable the integration of telemetric streams in connectivity health dashboards to help monitor battery life and residence of other critical sensor-enabled packages.

### Equation 5: Idempotent Processing Condition

#### Step 1: Define an operation

Let  $g$  be the consumer operation applied to state  $S$ .

Applying it once gives:

$$g(S)$$

Applying it twice gives:

$$g(g(S))$$

#### Step 2: Idempotent condition

By definition, an operation is idempotent if repeated application has the same effect as one application:


$$g(g(S)) = g(S)$$

#### Step 3: General repeated form

For any positive integer  $n$ ,

$$g^n(S) = g(S)$$

**5.2. Complex Event Processing and Real-Time Analytics**  
ACID transactional semantics limit query capabilities and,



because events are consumed and removed from the event store, it is not possible to conduct ad-hoc analyses or process events with different update frequencies. These constraints can be overcome through the use of an event-processing platform, which maintains continuously updated materialized views of the data (via atomic change data capture) that can then be consumed by analytic workloads with different requirements for data freshness and consistency.

In this setup, updates from an event stream are piped into the event processing (complex event-processing, CEP) system, where they can be aggregated, enriched with contextual data, correlated with other incoming streams, and persistently recorded. For on-demand queries requiring backfill of recently or regularly updated data, these views can be further served by a real-time analytics engine (such as Apache Druid), which performs optimizations specific to common patterns of dimensional querying. This supplementing approach is preferable if the operational costs of the event-processing system are high, or if the materialized views are not expected to be of sufficient quality to reliably support all analytic queries.

### 5.3. Data Lakes, Pipelines, and Streaming Platforms

Contemporary data streaming applications are supported by cloud-based data streaming platforms, typically based on the publish-subscribe pattern. These systems are capable of receiving, copying, and storing a large number of input streams for near-real-time delivery to multiple clients. Data streaming platforms support streaming applications by providing at least some of the following capabilities:

1. A unified framework that supports both standard data flow (request-and-reply) messaging and data streaming (publish-subscribe) without any programming effort.
2. High-throughput and low-latency data flow.
3. Storage of a large number of input streams with data in motion.

4. Ability to preserve message order within per-stream partitions.
5. In-stream query processing (within-stream analytics).
6. Audit capabilities for GDPR.
7. Automatic replication and fault tolerance capabilities of AWS services (such as Kinesis Streams, Kinesis Firehose, and S3).
8. Capability to connect existing data lakes (such as Amazon S3, Google Cloud Storage, or Hadoop) and existing "batch" analytics and reporting applications (such as Google BigQuery).
9. Ability to support built-in complex-event-processing (CEP) capabilities.

Storage systems designed specifically for streaming situations, often called streaming pipelines, are optimized for low-latency writing of data being received and high-throughput querying of data. Some support streaming query languages like SQL-on-Stream, SQL-like streaming languages, or any of the declarative programming languages supported by the existing ETL ecosystem.

## 6. Data Models and Semantics for Healthcare Events

Temporal clinical data can be structured as streams of events, providing valuable context when seeking to understand changes over time. Data semantics and events in the healthcare domain can align with the Fast Healthcare Interoperability Resources (FHIR) standard. Messages can be sent over a near real-time message bus for a publish-subscribe model to distribute to a number of interested parties. Provenance is critical for determining message source and audit trails for healthcare systems.


Much of the vast amount of healthcare data generated can be semantically represented using the FHIR standard. Such representation will enable communication and interoperability between many mainstream healthcare information systems, such as Electronic Health Records (EHRs), providing a sound basis for the monitoring of the complete healthcare journey of a patient.

Semantic alignment of events streamed over a near real-time messaging bus with FHIR provides a number of potential benefits. Semantically enriched data in a standard format allows rapid identification of a group of interest to whom messages can be distributed elsewhere in the system. Events sent to a streaming analytic processing capability can trigger alerts and provide input to complex event processing. A set of near-real-time event types captures the telemetry of the healthcare journey of a patient during acute care and after discharge from a care facility. Events are common change events not directly dominated by any specific healthcare standard.

## 6.1. FHIR and Health Data Semantics

Microservices rely on protocols and data formats that are understood by all adjacent components. In healthcare, organizations such as HL7 International define a suite of standards for the structure, semantics, and exchange of health data. FHIR is a family of common English domain models with machine-readable semantics, mostly expressed in the form of JSON objects. The use of URIs to designate resources, such as patient and provider identities, enables information to be stored independently in dedicated repositories for life-time management and integration into closed-loop systems e.g. at the time of medication fulfillment, eliminating the need for copy-and-paste of addressing data into orders, charts, etc. The FHIR resource definition includes a structure for tracing information about authorizations, accesses, modifications, and uses of the resource. Additional language modules use the well-formedness rules of FHIR team B (intensions).


In a federated (or decentralized) map Reduce architecture using file system and SQL commands, each involved organization has a local store for the full life cycle and governance of their health data and provides a wrapper process for each resource's content type listed in the uppercase level FHIR classification. Incoming data can be processed or ignored on demand, while outgoing data access respects the information-matching permission settings of FHIR, supporting the data-sharing capabilities of the HL7 FHIR Cloud. A translation process enables capturing in-event messages all occurrences of key resources across all organizations. The semantic mapping between the business process and the health event ontology serves in both detection and translation process.



**Fig 6: Health Data Integration with FHIR**

## 6.2. Temporal Alignment and Synchronization

Aligning events across multiple systems requires tightly synchronized clocks so that relative timestamps are semantically valid. When timestamps reflect different temporal domains, the placement of events along a single timeline loses meaning, making it impossible to understand their ordering and proximity. Despite the hard risks imposed on conventional clustering techniques, any approach relying on timestamp dissemination or network delay metrics can still shimmy around these issues by adopting directional intercommunication, a cosecant of red-blue networks in distributed systems. The Jason based protocol for reliable and time-stamped message delivery provides an effective solution but has yet to be examined in a highly dynamic context.



An alternative aligns events detected by multiple systems under a shared temporal domain in a non-intrusive manner. Snapshots of relevant state data is transferred to an authoring publisher by a pull-based data exposure pattern. Polling intervals on the subscriber designate epochs. Other event sources produce snapshots of detected events for the same epoch and publish these together with timestamped state snapshots under a publishing approach. By anchoring their timestamps at the beginning of the epoch, the order of relations in the generated newsletter maintains its derivation-semantics. While the approach can introduce nondeterminism, certain applications can adopt approximation-closed formats to trade accuracy for improved quality.

| Requirement     | Description                         |
|-----------------|-------------------------------------|
| Low Latency     | Minimal delay in processing         |
| High Throughput | Ability to handle large data volume |
| Reliability     | No data loss                        |
| Fault Tolerance | Recovery from failures              |
| Security        | Encryption, access control          |
| Privacy         | Data minimization, anonymization    |
| Compliance      | HIPAA, GDPR adherence               |
| Auditability    | Traceability of data access         |

**Table 3: Key Requirements for Near Real-Time Healthcare Streaming**

### 6.3. Provenance and Auditability

Data provenance refers to the ability to trace and understand the origin and evolution of data, answering questions such as: Who created it? When? How was it modified? Why? With what authority? Data provenance is a critical aspect in a trusted system, as the reliability of any derived inference

depends on the quality of the underlying data. In an EDA model designed for healthcare applications, data provenance should be closely observed for three reasons: (a) A device event notification concerns, in most cases, a person. An EDA system therefore needs a strategy to derive and bind to these events the associated individuals and/or groups of relevance. (b) With health and clinical information changing over time, it is desirable to maintain an up-to-date view that is consistent with the overall health information of the same individuals. (c) Key data in healthcare originate from processes governed by applicable clinical guidelines and protocols.

Audibility refers to the ability to reconstruct how a particular data set came about, including the completion of all operations along its life cycle, the relevant actors, and the reasons behind these actions. Audit logs not only assist in checking compliance with regulations, thereby assuring accountability, but also uncover patterns of legitimate data misuse that could be leveraged even to improve the system itself, such as by receptive learning. Audit logs thus need to be easily generated with minimal overhead while preventing unauthorized users from tampering with them.

## 7. Reliability, Consistency, and Fault Tolerance

Guaranteeing delivery semantics of exactly once—meaning an event is received by the subscriber just one time—implies the occurrence of every message both at the source and at the sink while preventing duplicates. Since network components are inherently unreliable, the more practical approach is to ensure least once delivery: the system attempts to deliver each message one or multiple times until it succeeds. Consequently, a receiver must tolerate duplicates. Accepting a leaky abstraction—in this case, an at-least-once delivery guarantee—calls for employing idempotency. Operations are designated idempotent when executing them multiple times has the same effect as executing them just once. Systems that do not allow such

redundancy must take care to write exactly-once, creating a bigger responsibility than before.


Handling duplicate streaming data also requires replays and the ability to recover from failures. These situations become evident in a visit to any streaming-platform user guide. As with any database, no matter how robustly fault-tolerant and scalable, streaming systems have failures. During a failure condition, messages are held in storage temporarily but can be accessed again when the consumer is back up and running. Specifically, stream processors use these mechanisms to implement exactly-once processing semantics in a system that does not support it natively. The concept of backpressure surfaces whenever processing becomes unable to keep up with the pace of incoming messages. Given high volumes, backpressure is an expected condition in a healthy streaming workload but should be a low-frequency event in an ideal pipeline. Nevertheless, it can create a problem if consuming applications are forced to drop messages. Implementing backpressure demands that every downstream activity denote its capacity to consume.

### 7.1. Exactly-Once vs At-Least-Once Delivery

Healthcare streams should be delivered reliably. In the context of a publish-subscribe mechanism, reliable stream delivery can be provided using redundant copies at source and destination. Within environments that support replication, such environments can additionally provide a form of delivery guarantee referred to as exactly-once delivery. This means that subscribers are guaranteed to receive an event/message exactly once. However, in an open environment involving multiple data producers and consumers, achieving exactly-once delivery is challenging and typically not feasible. More commonly, data messages are delivered to subscribers at least once, which means that subscribers may receive duplicate messages.

Handling duplicate messages requires that the consumers be idempotent with respect to the message processing. Idempotent means that applying the same operation more than once is equivalent to applying it once, i.e., it has no

additional effect beyond the first application. This property cannot be taken for granted in real-world applications. The duplicate delivery would need to be detected at the consumer side, and it requires maintaining some state information. In situations where the event/messages relate to an incoming data stream, duplicate detection is necessary temporally; that is, for a sliding time window over the consumer's recent history. The additional memory overhead associated with duplicate detection may vary based on the scalability aspects. Addressing this overhead in an efficient manner is an area of active research effort. Dedicated stream processing systems provide mechanisms/services to facilitate the development of scalable and resilient CEP applications, handling both strategies for fault tolerance: recovery and redundancy.



### 7.2. Idempotency and Replays

Exactly-once delivery is not always necessary or even desirable in healthcare applications. To maintain reliable workflows at high data volumes while providing loose data delivery and processing guarantees, idempotency and replay strategies can be employed. An operation is idempotent if executing it multiple times has the same effect as executing it only once. For example, ingesting a FHIR patient resource into a healthcare data warehouse can safely be performed at least once, provided that repeated ingestion of the same resource does not create inconsistencies due to the constraint

that a resource with the same id must be unique across the FHIR server and that previous versions remain accessible. Operations on "provisioned" entities or those on a logical control structure with client-side state reconciliation logic can also be designed to be idempotent. In addition to being idempotent, other operations within a data processing pipeline must also be designed to handle multiple executions while producing consistent results.

Replay is a strategy used to reconstruct the current state of an application or data source by reprocessing all, or a subset of, event data. An event data archive is a system that stores and provides access to event data for one or more use cases, such as debugging, responding to an industry requirement, recovering from a failure, validating an upgrade, or supporting an investigation. In a healthcare context, replaying events from a data lake enables organizations to view their past health information at a low cost, adhere to data controls, comply with regulatory requirements, and satisfy patient requests.

### 7.3. Resilience Patterns and Backpressure

To cope with unexpected conditions, automated resilience patterns must be applied. Such conditions include excessive unprocessed event accumulation in the processing stage and failure of all active processing instances. The messages of the topic-playback stage describe the occurrence of events processed later or after the time of publication. (These events are not sent to a playback topic when an exactly-once processing semantic is used.) Their consumption determines whether a fault is permanent or transient.


For cases with backpressure, it must be possible to defer the processing of events, thereby accepting an at-least-once delivery semantic for those events. During fault-free operations, the backpressure management prevents the processing of events that cannot be handled and delays their processing. When the conditions for handling such events are met (e.g., the trigger of fault detection is not in the list of available triggers), these events can be safely processed. The

sequence of events can be continued or resumed according to the fault nature.

## 8. Case Studies and Evaluation


Three Internet of Things (IoT) telemedicine systems demonstrate that current technology can support the requirements of near-real-time healthcare data streaming. An architecture for ICU patient monitoring sends a stream of physiological data to cloud-based analytics and alerting. Post-acute care for chronic obstructive pulmonary disease (COPD) uses health events from various systems and sources to drive recovery. A cardiac telemetry system transmits streams of telemetry events from wearable devices and medical implants.

Four characteristics distinguish the systems: streams of health data or telemetric telemetry data; near-real-time operation; processing, storage, or distribution performed by cloud systems; and extensive use of Internet health data models. Together, they validate the emerging architecture for a near-real-time health data stream.



### 8.1. Acute Care Monitoring

Close monitoring of patients in critical condition using telemetry tends to produce highly reliable event data at high frequencies. It is reasonable for these persistent telemetry streams to be integrated into the patient record as near-real-time data with low latency. Hospitals require an integrated



solution that addresses the underlying context (new patient admission, transfer, or discharge) while delivering an auditable history of persistent data streams from individual patients over time. The Healthcare Event Stream pattern describes how care data from devices integrated through an ICU hospital services team can be pushed to the patient record, improving data availability.

Telemetry data may be recorded continuously, or it may be pervasively sampled with ample operating margin. Changes are rare in comparison to the overall data rate. Arranging the individual patient records in a time-tree structure provides a convenient skeleton on which per-patient data can be coupled with supervisory context data.

Telemetry monitoring and control involve a surrounding set of point alerts in which sampled telemetry events cross entry protocols for either ICU or post-ICU telemetry telemetry. Point alert injections, unlike telemetry streams, play no role in temporally coupling data. The conventional supervision model respects GxE for detected critical conditions, while events exceeding clear thresholds open the door to GxE support for concurrent measurements in either arrival or severity.

## 8.2. Post-Acute and Population Health Streams

Each major care transition introduces a shift in the clinical attention gradient from a continuum of care to “transit,” where the potential for rapid change and clinical deterioration is heightened. Near-real-time monitoring of key health signals is essential to ensure these transitions do not become critical events. For post-acute care transitions, telemonitoring offers the potential for mentoring from an expert care team during recovery detailed in the American Telemedicine Association’s guidelines. Early trends across a federation of health systems also suggest that shared saved for transport and remote monitoring of sicker population is a good candidate for adoption. The risk of care being interrupted by re-hospitalization in patients with chronic diseases transitions must be carefully managed, both by a consortium of clinical experts and by the systems used to


support clinical decisions. Resilient, flexible architectures for near-real-time streaming should therefore be developed for monitoring these patients across both the time and service transition.

Streams of events and of patients at risk for re-hospitalization should therefore be detected and captured. Established clinical protocols or emerging clinical knowledge about these patients can produce key events for initiation of the monitoring service. Risk stratification for re-hospitalization can also be modelled as an ensemble of classifiers in a complex-event-processing engine. Risk becomes the monitored signal and the source of alarm-triggered actions. Candidate patients can enter a data shop where recent telemetry feeds, closed-loop therapy devices, care transition data managers, and an available care team are fused and checked for readiness. An advanced organizational technique such as a multi-agent system can be applied.

## 8.3. Telemetry from Medical Devices

Telemetry data from medical devices can supplement the continuous monitoring of patients in acute care settings as patients recover or transition into post-acute settings. Alerts and other clinically relevant information from devices continuously checking patients in Acute Care settings may be posted to an Event Broker for subscription by interested Services. In addition to these alerts, telemetry data generated by devices continually sampling a patient’s vital signs can also be periodically published to an Event Broker. Updates could occur with intervals bounded by a “least reflecting duration,” perhaps a “least specific” period or similar constraint. As a minimum requirement, device-compatible Event Message Types may implement structured present or availability elements so that a subscription event around Expected Patient Data Availability for Ventilators may be accurately checked.

Telemetry data from medical devices maintaining periodic contact with the Event Broker during Patient’s recovery would create a Continuous Patient Monitoring Data Stream during Near-Real-Time-Operation. The publications from



the devices also support state management and may be used to keep Devices and Monitoring Dashboards updated. Data from the recovery phase from telemetry devices can also be used for post-acute care transition or even for population health insights.

Telemetry data are also supported via the WebInterface of the Event Broker. During patient recovery, Event Messages can be fetched in Web Dashboard, or other Web Interfaces, of Event Message Queues/Topics managed by the Event Broker. Standard HTML Protocol-Plugin Interface of the Event Broker can also be used to read telemetry data for other purposes.

## 9. Conclusion

Research into near-real-time healthcare data streaming can be summarized as follows: an event-driven architecture closely aligned with evented systems leverages event-based and publish-subscribe patterns, allowing the auditing of data provenance, quality, and security; near-real-time streaming requires addressing latency, throughput, reliability, security, privacy, compliance, and interoperability considerations; supporting near-real-time streaming is straightforward with established architectural patterns such as streaming analytics, event sourcing, complex event processing, data pipelines, and data lakes; healthcare domains generate near-real-time data streams such as telemetry from medical devices, semi-automated acute-care monitoring, and semantic screens for post-acute and population-health care; well-structured semantic base models such as FHIR can describe and satisfy many healthcare event needs; and implementations can support exactly-once or at-least-once delivery guarantees while addressing fault tolerance and reliability requirements around idempotency and resilience.

Future work may also expand on these trends. The requirements imposed by new technologies or new societal expectations—such as remote monitoring, post-acute care of COVID-19 survivors, or the potential for psychometric monitoring through everyday devices—will create pressure

for early indicators of impending crises or pandemics, such as increasingly abnormal lung sounds in geographically linked cohorts. These sources justify the effort to create near-real-time pipelines and derive value from compressed sources of potentially overlapping data that quickly become stale. The quality offered by current methods and platforms for deep learning, integrated with event-processing infrastructures, suggests a strong aptitude for monitoring shifts in statistical distribution and spotting sudden emerging signs.

### 9.1. Emerging Trends

The digitization of healthcare data is generating unprecedented volumes of information that can be used for improved clinical decisions, population health monitoring, and the development of machine-learning models to enhance patient care. Traditionally, data has been collected and stored in data warehouses, batching, and time-slicing models. However, these methods introduce significant latency in getting the data visible downstream and also risk missing any rapid changes.

Emerging architecture patterns emphasize providing instantaneous access to event data and implementing near-real-time streaming from EHRs and other data sources. This has germinated a number of sub-disciplines, such as Complex Event Processing, Real-Time Analytics, and Predictive Analytics. The work analyzes the security, privacy, compliance, interoperability, and reliability requirements to develop these emergent patterns in a healthcare environment. Designed standards such as HL7 FHIR for healthcare data provide a formal semantic model enabling semantic interoperability. The inherent support for standardized semantic metadata makes EDA pattern implementations a real possibility and the development of patterns such as Publish-Subscribe, Event Sourcing, Data Lakes, and Streaming Platforms become practical.

# EUROPEAN ADVANCED JOURNAL FOR EMERGING TECHNOLOGIES

VOLUME: 04 ISSUE: 02

RECEIVED: APRIL 24

REVISED: MAY 12

ACCEPTED: MAY 26

PUBLISHED: JUNE 13



## 10. References

1. Abd-Alrazaq, A., Alajlani, M., Alhuwail, D., Schneider, J., Househ, M., Hamdi, M., & Shah, Z. (2023). Artificial intelligence in healthcare: A comprehensive review of its applications, challenges, and future directions. *Journal of Medical Systems*, 47(1), 6.
2. Agrawal, A., Choudhary, A., & Sharma, R. (2024). Real-time healthcare analytics using streaming data architectures. *IEEE Access*, 12, 27451–27469.
3. Inala, R., Sheelam, G. K., Aitha, A. R., Lakshmi, A. U., Nagabhyru, K. C., & Segireddy, A. R. (2026). Architecting Hybrid Data Products Using AI/ML and Agentic AI for Group Insurance and Retirement Solution Platforms with Advanced Data Governance. In *2026 IEEE International Conference on AI Engineering and Innovations (AIEI)* (pp. 1–6). IEEE. 2026 IEEE International Conference on AI Engineering and Innovations (AIEI). <https://doi.org/10.1109/aiei69164.2026.11497971>
4. Alowais, S. A., Alghamdi, S. S., Alsuehaby, N., Alqahtani, T., Alshaya, A. I., Almohareb, S. N., ... & Albekairy, A. M. (2023). Revolutionizing healthcare: The role of artificial intelligence in clinical practice. *BMC Medical Education*, 23, 689.
5. Ayaz, M., Pasha, M. F., Alzahrani, M. Y., Budiarto, R., & Stiawan, D. (2021). The Fast Healthcare Interoperability Resources (FHIR) standard: Systematic literature review of implementations, applications, challenges and opportunities. *JMIR Medical Informatics*, 9(7), e21929.
6. Kumar, M. V. K., Kolla, S. H., Pamisetty, V., Pandiri, L., Yandamuri, U. S., & Valiki, D. (2026). Enterprise-Scale Generative AI Agents for Secure and Governed Automation in Insurance and Public Financial Management. In *2026 International Conference on Computational Robotics, Testing and Engineering Evaluation (ICCRTEE)* (pp. 1–6). IEEE. 2026 International Conference on Computational Robotics, Testing and Engineering Evaluation (ICCRTEE). <https://doi.org/10.1109/iccrtee68719.2026.11566439>
7. Bashshur, R., Doarn, C. R., Frenk, J. M., Kvedar, J. C., & Woolliscroft, J. O. (2021). Telemedicine and the COVID-19 pandemic: Lessons for the future. *Telemedicine and e-Health*, 27(5), 571–573.
8. Bohr, A., & Memarzadeh, K. (2022). Artificial intelligence in healthcare. Academic Press.
9. Cheng, A. C., Duda, S. N., Taylor, R., Delacqua, F., Lewis, A. A., Bosler, T., Johnson, K. B., & Harris, P. A. (2021). REDCap on FHIR: Clinical data interoperability services. *Journal of Biomedical Informatics*, 121, 103871.
10. Radha, S., Gottimukkala, V. R. R., Thottara, S., Vandhana, K., & J, Gokulraj. (2025). Adaptive Video Streaming Over 5G Networks Using Deep Reinforcement Learning with Closed-Loop Feedback Mechanism for Bitrate Control. In *2025 International Conference on Communication, Computer, and Information Technology (IC3IT)* (pp. 1–6). IEEE. 2025 International Conference on Communication, Computer, and Information Technology (IC3IT). <https://doi.org/10.1109/ic3it66137.2025.11341184>
11. Dang, L. M., Piran, M. J., Han, D., Min, K., & Moon, H. (2021). A survey on Internet of Things and cloud computing for healthcare. *Electronics*, 10(7), 768.
12. Dash, S., Shakyawar, S. K., Sharma, M., & Kaushik, S. (2022). Big data in healthcare: Management, analysis, and future prospects. *Journal of Big Data*, 9(1), 54.
13. Hiremath, N. B., Kolla, S. K., Sunkara, R., G, S. Lal., & Sireesha, K. (2026). Adaptive and Intelligent Secure Multimedia Transmission for Dynamic Communication Systems. In *2026 2nd International Conference on Intelligent Systems and Computational Networks (ICISCN)* (pp. 1–8). IEEE. 2026 2nd International Conference on Intelligent Systems and Computational Networks (ICISCN). <https://doi.org/10.1109/iciscn67954.2026.11566143>
14. Ding, H., Wu, J., Zhao, Y., & Wang, X. (2024). Streaming analytics for intelligent healthcare

# EUROPEAN ADVANCED JOURNAL FOR EMERGING TECHNOLOGIES

VOLUME: 04 ISSUE: 02

RECEIVED: APRIL 24

REVISED: MAY 12

ACCEPTED: MAY 26

PUBLISHED: JUNE 13



- monitoring systems. *Future Generation Computer Systems*, 154, 352–366.
15. El-Sappagh, S., Ali, F., Hendawi, A., Jang, J. H., & Kwak, K. S. (2022). Real-time healthcare monitoring using Internet of Medical Things and cloud computing. *Sensors*, 22(4), 1450.
  16. Esteva, A., Robicquet, A., Ramsundar, B., Kuleshov, V., DePristo, M., Chou, K., ... & Dean, J. (2021). A guide to deep learning in healthcare. *Nature Medicine*, 27(1), 24–29.
  17. Baladari, V., Nagubandi, A. R., Charan Teja Tadi, S. R. C., Selvi, A. T., Sreedevi, V., & Nithya, M. (2026). Predictive AI Model for Financial Risk Assessment in Dynamic Market Environments. In 2026 International Conference on Communication, Computing and Emerging Technologies (IC3ET) (pp. 748–753). IEEE. 2026 International Conference on Communication, Computing and Emerging Technologies (IC3ET). <https://doi.org/10.1109/ic3et64989.2026.11467452>
  18. Garcia, S., Kumar, P., & Singh, V. (2025). Adaptive healthcare data streaming for smart hospitals. *IEEE Internet of Things Journal*, 12(2), 1950–1964.
  19. Hossain, M. S., Muhammad, G., & Guizani, N. (2022). Cloud-assisted industrial Internet of Things for healthcare applications. *IEEE Network*, 36(2), 136–143.
  20. Islam, S. M. R., Kwak, D., Kabir, M. H., Hossain, M., & Kwak, K. S. (2021). The Internet of Things for health care: A comprehensive survey. *IEEE Access*, 9, 67884–67910.
  21. Kumar, S. S., Garapati, R. S., Segireddy, A. R., Kalisetty, S., Inala, R., & Nagabhyru, K. C. (2026). Hybrid Deep Neural Network–DevOps Pipeline Optimization for Risk Prediction in Cloud-Native Workers' Compensation Platforms. In 2026 IEEE International Conference for Convergence in Computing Technology (I3CTCON) (pp. 1–6). IEEE. 2026 IEEE International Conference for Convergence in Computing Technology (I3CTCON). <https://doi.org/10.1109/i3ctcon68242.2026.11507219>
  22. Kamel Boulos, M. N., Peng, G., & VoPham, T. (2023). An overview of healthcare data interoperability and digital transformation. *International Journal of Health Geographics*, 22(1), 17.
  23. Khan, R. A., Jawaaid, M., Khan, A. R., & Sajjad, M. (2024). Artificial intelligence for real-time patient monitoring: A systematic review. *Healthcare*, 12(3), 401.
  24. Li, X., Zhang, Y., Wang, J., & Chen, L. (2025). Event-driven architectures for responsive healthcare information systems. *Future Internet*, 17(1), 25.
  25. Bedi, B., Yandamuri, U. S., Kummari, D. N., Nagubandi, A. R., Amistapuram, K., & D, A. (2026). AI-Driven Sentiment and Behavior Analysis for Sustainable Business Growth. In 2026 International Conference on Emerging Research in Smart Electronics and Machine Informatics (ECMI) (pp. 1–11). IEEE. 2026 International Conference on Emerging Research in Smart Electronics and Machine Informatics (ECMI). <https://doi.org/10.1109/ecmi68341.2026.11603189>
  26. Luo, J., Wu, M., Gopukumar, D., & Zhao, Y. (2022). Big data application in biomedical research and healthcare: A literature review. *Biomedical Informatics Insights*, 14, 1–18.
  27. Meskó, B., Drobni, Z., Bényei, É., Gergely, B., & Györfy, Z. (2022). Digital health technologies and the future of healthcare. *NPJ Digital Medicine*, 5(1), 89.
  28. Reddy, M. S. R. L., Sunitha, T., Kanchana, K., Nagubandi, A. R., Segireddy, A. R., & Bhavanam, S. N. (2026). AI-Enhanced Blockchain Consensus Mechanisms for Secure Transaction Validation. In 2026 International Conference on Emerging Research in Smart Electronics and Machine Informatics (ECMI) (pp. 1–11). IEEE. 2026 International Conference on Emerging Research in Smart Electronics and Machine Informatics (ECMI). <https://doi.org/10.1109/ecmi68341.2026.11602724>
  29. Nair, M., & Kaur, P. (2023). Secure healthcare data exchange using blockchain and FHIR

# EUROPEAN ADVANCED JOURNAL FOR EMERGING TECHNOLOGIES

VOLUME: 04 ISSUE: 02

RECEIVED: APRIL 24

REVISED: MAY 12

ACCEPTED: MAY 26

PUBLISHED: JUNE 13



- interoperability. *Journal of Network and Computer Applications*, 218, 103658.
30. Rojas, E., Munoz-Gama, J., Sepúlveda, M., & Capurro, D. (2021). Process mining in healthcare: A literature review. *Journal of Biomedical Informatics*, 121, 103891.
  31. Shah, S. A., Khan, M. A., Abbas, S., Athar, A., Khan, B. S., & Aslam, A. (2024). Intelligent healthcare systems using IoT and edge computing: Recent advances and challenges. *Sensors*, 24(2), 781.
  32. Subramanian, N., & Kolla, T. Equity in Healthcare Access Policy Lessons from Universal Coverage Models.
  33. Topol, E. (2023). High-performance medicine: The convergence of human and artificial intelligence. *Nature Medicine*, 29(1), 44–56.
  34. Wang, L., Zhang, H., Liu, Y., & Chen, Z. (2025). Responsive healthcare data pipelines for clinical decision support in distributed environments. *Journal of Biomedical Informatics*, 157, 104623.
  35. Wang, S., Wang, Y., Zhang, X., & Li, J. (2023). Edge intelligence for real-time healthcare monitoring: A survey. *IEEE Internet of Things Journal*, 10(14), 12136–12158.
  36. Sudha Rani, P. R., Amistapuram, K., Pamisetty, V., Singireddy, S., Kummari, D. N., & Sheelam, G. K. (2025). Hybrid Knowledge Graph–Deep Learning Framework for Automated Exception Handling and Investigation in Complex Insurance Claims. In 2025 IEEE 3rd Global Conference on Wireless Computing and Networking (GCWCN) (pp. 1–6). IEEE. 2025 IEEE 3rd Global Conference on Wireless Computing and Networking (GCWCN).  
<https://doi.org/10.1109/gcwcnc66157.2025.11448301>
  37. Wu, H., Li, Y., Zhang, Z., & Sun, L. (2022). Healthcare Internet of Things: Architecture, enabling technologies, security and privacy. *IEEE Communications Surveys & Tutorials*, 24(2), 1185–1224.
  38. Xu, B., Li, H., Zhang, J., & Wang, Y. (2023). Intelligent health data integration using cloud-edge computing. *Future Generation Computer Systems*, 141, 228–241.
  39. Yang, G., He, Y., & Li, Z. (2024). Digital twins for smart healthcare: Recent advances and future opportunities. *IEEE Reviews in Biomedical Engineering*, 17, 75–92.
  40. Kumar, S. S., Garapati, R. S., Segireddy, A. R., Kalisetty, S., Inala, R., & Nagabhyru, K. C. (2026, March). Hybrid Deep Neural Network–DevOps Pipeline Optimization for Risk Prediction in Cloud-Native Workers’ Compensation Platforms. In 2026 IEEE International Conference for Convergence in Computing Technology (I3CTCON) (pp. 1–6). IEEE.
  41. Zhang, P., Walker, M. A., & Patel, V. L. (2021). Data interoperability in healthcare information systems: Current status and future perspectives. *Journal of Biomedical Informatics*, 122, 103906.
  42. Zhang, Y., Chen, X., & Wang, L. (2025). Scalable event-driven architectures for next-generation healthcare information systems. *Future Generation Computer Systems*, 156, 182–197.
  43. Zhao, J., Liu, X., Wang, T., & Huang, Z. (2023). Artificial intelligence-enabled healthcare data management: A review. *Information Fusion*, 92, 86–101.
  44. Thutari, R. T., Garapati, R. S., B M, Manjula., R K, Supriya., & M, Senbagan. (2025). Adaptive Access Control and Authentication Management for IoT Using Attention-GRU and Reinforcement Learning. In 2025 2nd International Conference on Software, Systems and Information Technology (SSITCON) (pp. 1–6). IEEE. 2025 2nd International Conference on Software, Systems and Information Technology (SSITCON).  
<https://doi.org/10.1109/ssitcon66133.2025.11342003>
  45. Albahri, O. S., Albahri, A. S., Mohammed, K. I., Zaidan, A. A., Zaidan, B. B., Hashim, M., & Alsalem, M. A. (2022). Systematic review of real-time remote health monitoring systems. *Journal of Medical Systems*, 46(9), 58.
  46. Benis, A., Tamburis, O., Chronaki, C., & Moen, A. (2021). One digital health: A unified framework for

# EUROPEAN ADVANCED JOURNAL FOR EMERGING TECHNOLOGIES

VOLUME: 04 ISSUE: 02

RECEIVED: APRIL 24

REVISED: MAY 12

ACCEPTED: MAY 26

PUBLISHED: JUNE 13



- future health ecosystems. *Journal of Medical Internet Research*, 23(2), e22189.
47. Chen, M., Ma, Y., Song, J., Lai, C. F., & Hu, B. (2021). Smart healthcare: Applications, challenges and future directions. *Journal of Network and Computer Applications*, 184, 103035.
  48. Padma, G., Reddy, V. A. R., KA, S. D., Buvaneswari, B., & Kumar, K. S. (2026, April). Privacy-Preserved Face Recognition Biometric Authentication Using FaceNet and Zero-Knowledge Proofs for Secure, Access Control on Decentralized Blockchain Networks. In *2026 2nd International Conference on Intelligent Systems and Computational Networks (ICISCN)* (pp. 1-8). IEEE.
  49. Dinh-Le, C., Chuang, R., Chokshi, S., & Mann, K. (2022). Wearable health technology and electronic health record integration: A systematic review. *Journal of Medical Systems*, 46(4), 24.
  50. Giansanti, D. (2023). Artificial intelligence in digital health: Current challenges and future perspectives. *Healthcare*, 11(12), 1712.
  51. Haleem, A., Javaid, M., Singh, R. P., & Suman, R. (2022). Telemedicine for healthcare: Capabilities, features, barriers, and applications. *Sensors International*, 3, 100117.
  52. Krishnan, M., Aitha, A. R., Amistapuram, K., Nandan, B. P., Kaulwar, P. K., & Singireddy, J. (2025). Human-in-the-Loop Hybrid Neuro-Symbolic AI Model for Reliable Data Engineering in High-Stakes Industrial Systems. In *2025 IEEE 3rd Global Conference on Wireless Computing and Networking (GCWCN)* (pp. 1–7). IEEE. 2025 IEEE 3rd Global Conference on Wireless Computing and Networking (GCWCN).  
<https://doi.org/10.1109/gcwcN66157.2025.11448516>
  53. He, J., Baxter, S. L., Xu, J., Xu, J., Zhou, X., & Zhang, K. (2021). The practical implementation of artificial intelligence technologies in medicine. *Nature Medicine*, 27(1), 31–38.
  54. Islam, M. M., Rahaman, A., & Islam, M. R. (2022). Development of smart healthcare monitoring systems in IoT environment: A review. *Journal of King Saud University—Computer and Information Sciences*, 34(10), 8836–8853.
  55. Nagabhyru, K. C., & Priya, B. D. (2026, July). Physical-Unclonable-Function-Based Secure and Anonymous User Authentication for Smart Homes. In *Signal Processing, Telecommunication & Embedded Systems: Automation and Sustainability Applications: Proceedings of Tenth International Conference on Microelectronics Electromagnetics and Telecommunications (ICMEET 2025)*, Volume 4 (Vol. 4, p. 367). Springer Nature.
  56. Keesara, S., Jonas, A., & Schulman, K. (2021). Covid-19 and health care's digital revolution. *New England Journal of Medicine*, 384(23), e82.
  57. Khan, N., McClean, S., Nugent, C., Cleland, I., Zhang, S., & Scotney, B. (2023). Streaming analytics for wearable healthcare systems: A systematic review. *IEEE Access*, 11, 84526–84549.
  58. Garapati, R. S., Aitha, A. R., & Singireddy, S. (2026). Secure Cloud Architecture for Privacy-Preserving Machine Learning on Electronic Health Records with Web-Based Analytics Tools. In *International Conference on Intelligent Human Computer Interaction* (pp. 407–417). Springer, Cham.
  59. Kumar, A., Gupta, P. K., & Srivastava, A. (2024). Secure cloud-native healthcare platforms using container orchestration. *Journal of Systems Architecture*, 149, 103078.
  60. Li, F., Yu, H., Yang, J., & Zhao, W. (2025). Event stream processing for intelligent clinical decision support. *Journal of Biomedical Informatics*, 158, 104664.
  61. Martínez-Pérez, B., de la Torre-Díez, I., & López-Coronado, M. (2021). Mobile health applications for healthcare delivery: Current status and future perspectives. *Journal of Medical Systems*, 45(8), 74.
  62. Bhavani, B. D., SR, S., Loganathan, R., & Nagaraj, S. (2026, April). Evolutionary Gravitational Neocognitron Neural Network, Snow Leopard Optimization and Deep Graph Reinforcement Learning for Routing Protocol in WSN. In *2026 2nd*

# EUROPEAN ADVANCED JOURNAL FOR EMERGING TECHNOLOGIES

VOLUME: 04 ISSUE: 02

RECEIVED: APRIL 24

REVISED: MAY 12

ACCEPTED: MAY 26

PUBLISHED: JUNE 13



- International Conference on Intelligent Systems and Computational Networks (ICISCN) (pp. 1-8). IEEE.
63. Nguyen, D. C., Pathirana, P. N., Ding, M., & Seneviratne, A. (2021). Blockchain for secure electronic health records and healthcare systems: A survey. *IEEE Access*, 9, 101054–101079.
  64. Reddy, S., Allan, S., Coghlan, S., & Cooper, P. (2023). A governance model for the application of AI in healthcare. *Journal of the American Medical Informatics Association*, 30(3), 497–504.
  65. RK, S., Mangala, N., Priyanka, R., Sait, R. A. M., & Selvam, P. P. (2026, April). Privacy Preserving Analytics for Intelligent in Internet of Things System in Health Care Using Graph Neural Network with Latent Graph Inference Technique. In 2026 2nd International Conference on Intelligent Systems and Computational Networks (ICISCN) (pp. 1-6). IEEE.
  66. Shuaib, M., Samad, A., Alam, S., & Siddiqui, S. T. (2022). Why blockchain for healthcare? Electronic health record security and interoperability. *Sensors*, 22(13), 4794.
  67. Verma, P., Sood, S. K., & Kalra, S. (2024). Intelligent edge-cloud framework for real-time patient monitoring. *IEEE Transactions on Network and Service Management*, 21(2), 1735–1749.
  68. Zhang, H., Liu, Y., Chen, Q., & Wang, X. (2025). Real-time distributed healthcare data processing using event-driven microservices. *IEEE Access*, 13, 11544–11561.
  69. Paramasivam, R., Reddy, S. S., Reddy, V. A. R., Sandra, K., & Narayana, M. V. S. (2026, April). JellyFish Search Optimization with Arctic Puffin Optimization Algorithm for Efficient Routing Based on the Software Defined Communication Network. In 2026 2nd International Conference on Intelligent Systems and Computational Networks (ICISCN) (pp. 1-6). IEEE.
  70. Greer, M. L., Williams, T. B., Garza, M. Y., et al. (2023). Measuring the coverage of the HL7® FHIR® standard in supporting data acquisition for three public health registries. *Journal of Medical Systems*, 47(1), 112.
  71. Haleem, A., Javaid, M., Singh, R. P., & Suman, R. (2023). Digital healthcare technologies and their applications in healthcare systems. *Sensors International*, 4, 100238.
  72. Rao, C. S., Bandi, V. D. V. K., Brahmandam, L. M. K., Gantikota, S., & Kannan, K. (2026, April). Topology-Aware Hypergraph Neural Network Framework for Intelligent Decision Support Systems in Network Operations. In 2026 2nd International Conference on Intelligent Systems and Computational Networks (ICISCN) (pp. 1-7). IEEE.
  73. He, J., Baxter, S. L., Xu, J., Xu, J., Zhou, X., & Zhang, K. (2021). The practical implementation of artificial intelligence technologies in medicine. *Nature Medicine*, 27(1), 31–38.
  74. Keesara, S., Jonas, A., & Schulman, K. (2021). Covid-19 and health care's digital revolution. *New England Journal of Medicine*, 384(23), e82.
  75. Manikandan, S., Singh, G. P., Gottimukkala, V. R. R., Davuluri, P. N., Sadagopan, R., & Kumar, T. V. (2026, March). Human-in-the-Loop Automation Patterns for Financial Services Using Camunda+ Modern Java UIs. In 2026 IEEE International Conference for Convergence in Computing Technology (I3CTCON) (pp. 1-6). IEEE.
  76. Nitulescu, A., & Stoicu-Tivadar, L. (2024). Data standardization in the medical field through FHIR and FAIR implementation: A systematic review. *Studies in Health Technology and Informatics*, 316, 560–564.
  77. Pavão, J., Bastardo, R., Santos, M., & Rocha, N. P. (2023). The Fast Healthcare Interoperability Resources (FHIR) standard and homecare: A scoping review. *Procedia Computer Science*, 219, 1249–1256.
  78. Nagabhyru, K. C., Inala, R., Jayakumar, S., & Raj, S. R. (2026). Communication Resilience in Smart Grid Nans: Challenges, Solutions and Future Outlook. In *International Conference on Microelectronics, Electromagnetics and Telecommunication* (pp. 318-329). Springer, Cham.
  79. Pimenta, N., Chaves, A., Sousa, R., Abelha, A., & Peixoto, H. (2023). Interoperability of clinical data

# EUROPEAN ADVANCED JOURNAL FOR EMERGING TECHNOLOGIES

VOLUME: 04 ISSUE: 02

RECEIVED: APRIL 24

REVISED: MAY 12

ACCEPTED: MAY 26

PUBLISHED: JUNE 13



- through FHIR: A review. *Procedia Computer Science*, 220, 856–861.
80. Reddy, S., Allan, S., Coghlan, S., & Cooper, P. (2023). A governance model for the application of AI in healthcare. *Journal of the American Medical Informatics Association*, 30(3), 497–504.
  81. Gopinath, A., Davuluri, P. N., Kumar, U. B., MP, S., & Yamini, G. (2026, April). Communication Aware Malware Detection Using Network Flow Bytecode Fusion With Adaptive Focal Optimization. In 2026 2nd International Conference on Intelligent Systems and Computational Networks (ICISCN) (pp. 1-8). IEEE.
  82. Tabari, P., Costagliola, G., De Rosa, M., & Boeker, M. (2024). State-of-the-art Fast Healthcare Interoperability Resources (FHIR)-based data model and structure implementations: Systematic scoping review. *JMIR Medical Informatics*, 12, e58445.
  83. Torab-Miandoab, A., Samad-Soltani, T., Jodati, A., & Rezaei-Hachesu, P. (2023). Interoperability of heterogeneous health information systems: A systematic literature review. *BMC Medical Informatics and Decision Making*, 23, 18.
  84. Inala, R., Sheelam, G. K., Aitha, A. R., Lakshmi, A. U., Nagabhyru, K. C., & Segireddy, A. R. (2026, March). Architecting Hybrid Data Products Using AI/ML and Agentic AI for Group Insurance and Retirement Solution Platforms with Advanced Data Governance. In 2026 IEEE International Conference on AI Engineering and Innovations (AIEI) (pp. 1-6). IEEE.
  85. Gour, S., Peake, A., Tong, C., Churm, J., Ahmad, B., Pournik, O., & Arvanitis, T. N. (2024). Advancing healthcare through interoperability: Implementing scalable solutions for patient data integration. *Studies in Health Technology and Informatics*, 316, 242–246.
  86. Benis, A., Tamburis, O., Chronaki, C., & Moen, A. (2021). One Digital Health: A unified framework for future health ecosystems. *Journal of Medical Internet Research*, 23(2), e22189.
  87. Nayak, P., Loganathan, R., & Jayaraj, V. (2026, April). Scale-Aware Dilated Lightweight Convolutional Network Improving Solar Panel Defect Classification through Efficient Electroluminescent Image Analysis Techniques. In 2026 2nd International Conference on Intelligent Systems and Computational Networks (ICISCN) (pp. 1-7). IEEE.
  88. Giansanti, D. (2023). Artificial intelligence in digital health: Current challenges and future perspectives. *Healthcare*, 11(12), 1712.
  89. Alowais, S. A., Alghamdi, S. S., Alsuehaby, N., et al. (2023). Revolutionizing healthcare: The role of artificial intelligence in clinical practice. *BMC Medical Education*, 23, 689.
  90. Kolla, T. (2026). Blockchain for Health Records Management Policy, Legal, and Technical Perspectives. *Journal of Health Politics, Policy and Law*.
  91. Ayaz, M., Pasha, M. F., Alzahrani, M. Y., Budiarto, R., & Stiawan, D. (2021). The Fast Healthcare Interoperability Resources (FHIR) standard: Systematic literature review of implementations, applications, challenges, and opportunities. *JMIR Medical Informatics*, 9(7), e21929.
  92. Cheng, A. C., Duda, S. N., Taylor, R., Delacqua, F., Lewis, A. A., Bosler, T., Johnson, K. B., & Harris, P. A. (2021). REDCap on FHIR: Clinical data interoperability services. *Journal of Biomedical Informatics*, 121, 103871.
  93. Recharla, M., Garapati, R. S., Bandi, V. D. V. K., Yandamuri, U. S., & Mangalampalli, B. M. (2026, March). Generative AI-Enhanced Data Engineering Pipelines for Predictive Biomarker Discovery in Alzheimer's and Kidney Disease. In 2026 IEEE International Conference on AI Engineering and Innovations (AIEI) (pp. 1-5). IEEE.
  94. Dinh-Le, C., Chuang, R., Chokshi, S., & Mann, K. (2022). Wearable health technology and electronic health record integration: A systematic review. *Journal of Medical Systems*, 46(4), 24.
  95. El-Sappagh, S., Ali, F., Hendawi, A., Jang, J. H., & Kwak, K. S. (2022). Real-time healthcare monitoring using Internet of Medical Things and cloud computing. *Sensors*, 22(4), 1450.

# EUROPEAN ADVANCED JOURNAL FOR EMERGING TECHNOLOGIES

VOLUME: 04 ISSUE: 02

RECEIVED: APRIL 24

REVISED: MAY 12

ACCEPTED: MAY 26

PUBLISHED: JUNE 13



96. Islam, M. M., Rahaman, A., & Islam, M. R. (2022). Development of smart healthcare monitoring systems in IoT environment: A review. *Journal of King Saud University – Computer and Information Sciences*, 34(10), 8836–8853.
97. Dawadi, D., Yandamuri, U. S., Kumar, S., Stalin, J. A., & Naveenkumar, R. (2026, February). Privacy-Aware Edge-Based Intelligent Video Analytics for Scalable Crowd Management in Smart Cities. In 2026 3rd International Conference on Integrated Intelligence and Communication Systems (ICIICS) (pp. 1-7). IEEE.
98. Luo, J., Wu, M., Gopukumar, D., & Zhao, Y. (2022). Big data application in biomedical research and healthcare: A literature review. *Biomedical Informatics Insights*, 14, 1–18.
99. Meskó, B., Drobni, Z., Bényei, É., Gergely, B., & Györfi, Z. (2022). Digital health technologies and the future of healthcare. *NPJ Digital Medicine*, 5(1), 89.
100. Nguyen, D. C., Pathirana, P. N., Ding, M., & Seneviratne, A. (2021). Blockchain for secure electronic health records and healthcare systems: A survey. *IEEE Access*, 9, 101054–101079.
101. Singh, D., Peruri, S. V. S., Mangala, N., Gurram, R. T., & Sai, D. H. (2026, April). Privacy-Aware Data Sharing Using Improved Distributed Ledger Technologies in Intelligent Communication Networks. In 2026 2nd International Conference on Intelligent Systems and Computational Networks (ICISCN) (pp. 1-6). IEEE.
102. Shuaib, M., Samad, A., Alam, S., & Siddiqui, S. T. (2022). Why blockchain for healthcare? Electronic health record security and interoperability. *Sensors*, 22(13), 4794.
103. Wang, S., Wang, Y., Zhang, X., & Li, J. (2023). Edge intelligence for real-time healthcare monitoring: A survey. *IEEE Internet of Things Journal*, 10(14), 12136–12158.
104. Kuppuswami, R., Yandamuri, U. S., Bhatt, M., & Patel, M. (2026, February). A Cognitive Clustering Framework for Wireless Sensor Networks Using Quantum Machine Learning. In 2026 3rd International Conference on Integrated Intelligence and Communication Systems (ICIICS) (pp. 1-6). IEEE.
105. Zhao, J., Liu, X., Wang, T., & Huang, Z. (2023). Artificial intelligence-enabled healthcare data management: A review. *Information Fusion*, 92, 86–101.
106. Bastola, N. D., Tcheng, J. E., Schlossman, D. M., & Windle, J. R. (2025). Framework for improving patient safety: Reference model for FHIR-enabled, patient-centric home medication list management and medication reconciliation. *Applied Clinical Informatics*, 16(4), 1136–1145.
107. Kummari, D. N., Aitha, A. R., Kumar, M. V. K., Pandiri, L., Gottimukkala, V. R. R., & Nagubandi, A. R. (2026, March). Real-Time AI-Driven Audit and Compliance Framework for Smart Manufacturing Financial Workflow. In 2026 IEEE International Conference on AI Engineering and Innovations (AIEI) (pp. 1-5). IEEE.
108. Cheng, A. C., Shyr, C., Lewis, A., Delacqua, F., Bosler, T., Banasiewicz, M. K., & Harris, P. A. (2025). Opportunities, barriers, and remedies for implementing REDCap integration with electronic health records via Fast Healthcare Interoperability Resources (FHIR). *JAMIA Open*, 8(5).
109. Leaphart, D., Richesson, R. L., Correa, K. H., Ahmed, A., Boxwala, A. A., Desai, P. J., & Dullabh, P. M. (2025). Making patient apps interoperable with the health IT ecosystem. *Agency for Healthcare Research and Quality*.
110. Jyoshna, K., Peddi, R. K., & Punitha, S. (2026, April). Spatio-Temporal Graph Neural Network with Global Spatio-Temporal Network for the Traffic Flow Prediction. In 2026 2nd International Conference on Intelligent Systems and Computational Networks (ICISCN) (pp. 1-6). IEEE.
111. Löbe, M., Bartow, M., Bialke, M., Gatrio, M., Gulden, C., Sass, J., et al. (2025). Supporting interoperability and transparency of medical research studies with HL7 FHIR. In *MEDINFO 2025* –

# EUROPEAN ADVANCED JOURNAL FOR EMERGING TECHNOLOGIES

VOLUME: 04 ISSUE: 02

RECEIVED: APRIL 24

REVISED: MAY 12

ACCEPTED: MAY 26

PUBLISHED: JUNE 13



- Healthcare Smart × Medicine Deep (pp. 43–47). IOS Press.
112. Torab-Miandoab, A., Samad-Soltani, T., Jodati, A., & Rezaei-Hachesu, P. (2023). Interoperability of heterogeneous health information systems: A systematic literature review. *BMC Medical Informatics and Decision Making*, 23, 18.
  113. Tabari, P., Costagliola, G., De Rosa, M., & Boeker, M. (2024). State-of-the-art Fast Healthcare Interoperability Resources (FHIR)-based data model and structure implementations: Systematic scoping review. *JMIR Medical Informatics*, 12, e58445.
  114. Hiremath, N. B., Kolla, S. K., Sunkara, R., & Sireesha, K. (2026, April). Adaptive and Intelligent Secure Multimedia Transmission for Dynamic Communication Systems. In *2026 2nd International Conference on Intelligent Systems and Computational Networks (ICISCN)* (pp. 1-8). IEEE.
  115. Nitulescu, A., & Stoicu-Tivadar, L. (2024). Data standardization in the medical field through FHIR and FAIR implementation: A systematic review. *Studies in Health Technology and Informatics*, 316, 560–564.
  116. Pavão, J., Bastardo, R., Santos, M., & Rocha, N. P. (2023). The Fast Healthcare Interoperability Resources (FHIR) standard and homecare: A scoping review. *Procedia Computer Science*, 219, 1249–1256.
  117. Pimenta, N., Chaves, A., Sousa, R., Abelha, A., & Peixoto, H. (2023). Interoperability of clinical data through FHIR: A review. *Procedia Computer Science*, 220, 856–861.
  118. Davuluri, P. N. AI-Augmented Sanctions Screening: Enhancing Accuracy and Latency in Real Time Compliance Systems.
  119. Greer, M. L., Williams, T. B., Garza, M. Y., et al. (2023). Measuring the coverage of the HL7 FHIR standard in supporting data acquisition for public health registries. *Journal of Medical Systems*, 47(1).
  120. Benis, A., Tamburis, O., Chronaki, C., & Moen, A. (2021). One Digital Health: A unified framework for future health ecosystems. *Journal of Medical Internet Research*, 23(2), e22189.
  121. Giansanti, D. (2023). Artificial intelligence in digital health: Current challenges and future perspectives. *Healthcare*, 11(12), 1712.
  122. Alowais, S. A., Alghamdi, S. S., Alsuehany, N., et al. (2023). Revolutionizing healthcare: The role of artificial intelligence in clinical practice. *BMC Medical Education*, 23, 689.
  123. Cheng, A. C., Duda, S. N., Taylor, R., Delacqua, F., Lewis, A. A., Bosler, T., Johnson, K. B., & Harris, P. A. (2021). REDCap on FHIR: Clinical data interoperability services. *Journal of Biomedical Informatics*, 121, 103871.
  124. Ayaz, M., Pasha, M. F., Alzahrani, M. Y., Budiarto, R., & Stiawan, D. (2021). The Fast Healthcare Interoperability Resources (FHIR) standard: Systematic literature review of implementations, applications, challenges, and opportunities. *JMIR Medical Informatics*, 9(7), e21929.
  125. Reddy, C. S., Kolla, S. H., Kumar, R. D., & Koteswararao, O. V. (2026, April). A Temporal Attention Transformer Framework Based Air Quality Monitoring and Forecasting in Urban Environments. In *2026 2nd International Conference on Intelligent Systems and Computational Networks (ICISCN)* (pp. 1-6). IEEE.
  126. Dinh-Le, C., Chuang, R., Chokshi, S., & Mann, K. (2022). Wearable health technology and electronic health record integration: A systematic review. *Journal of Medical Systems*, 46(4), 24.
  127. El-Sappagh, S., Ali, F., Hendawi, A., Jang, J. H., & Kwak, K. S. (2022). Real-time healthcare monitoring using Internet of Medical Things and cloud computing. *Sensors*, 22(4), 1450.
  128. Islam, M. M., Rahaman, A., & Islam, M. R. (2022). Development of smart healthcare monitoring systems in IoT environment: A review. *Journal of King Saud University – Computer and Information Sciences*, 34(10), 8836–8853.
  129. Luo, J., Wu, M., Gopukumar, D., & Zhao, Y. (2022). Big data application in biomedical research and healthcare: A literature review. *Biomedical Informatics Insights*, 14, 1–18.

130. Meskó, B., Drobni, Z., Bényei, É., Gergely, B., & Györfy, Z. (2022). Digital health technologies and the future of healthcare. *NPJ Digital Medicine*, 5(1), 89.
131. Nguyen, D. C., Pathirana, P. N., Ding, M., & Seneviratne, A. (2021). Blockchain for secure electronic health records and healthcare systems: A survey. *IEEE Access*, 9, 101054–101079.
132. Shuaib, M., Samad, A., Alam, S., & Siddiqui, S. T. (2022). Why blockchain for healthcare? Electronic health record security and interoperability. *Sensors*, 22(13), 4794.
133. Reddy, S., Allan, S., Coghlan, S., & Cooper, P. (2023). A governance model for the application of AI in healthcare. *Journal of the American Medical Informatics Association*, 30(3), 497–504.
134. Haleem, A., Javaid, M., Singh, R. P., & Suman, R. (2023). Digital healthcare technologies and their applications in healthcare systems. *Sensors International*, 4, 100238.
135. He, J., Baxter, S. L., Xu, J., Xu, J., Zhou, X., & Zhang, K. (2021). The practical implementation of artificial intelligence technologies in medicine. *Nature Medicine*, 27(1), 31–38.